Beyond the Constraints: Unleashing the Potential of Low-Cost Robots and Standardizing Robotics Performance Assessment

Boston University

CS599N1: Robot Brains! Designing Computing Systems for Robotics

Oct. 31, 2023

Jason Jabbour

CS PhD Student @ Harvard University

Roadmap

- Part 1 - **Tiny Robot Learning** + Class Discussion (9:35am)
- Part 2 - **RobotPerf** + Class Discussion (10:10am)

About Me

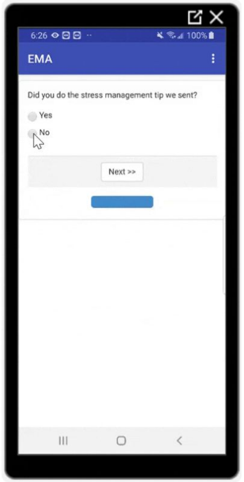

6:26
100%
EMA
Did you do the stress management tip we sent?
Yes
No
Next >>

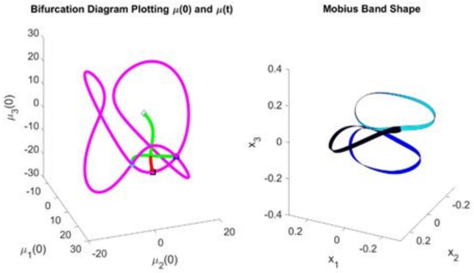

Bifurcation Diagram Plotting $\mu(0)$ and $\mu(t)$
$\mu_3(0)$
$\mu_1(0)$
$\mu_2(0)$
Mobius Band Shape
x_3
x_1
x_2

Network Traffic
Position
Orientation
Velocity
Acceleration
Wind Speed
Sensors
Subscribers
Navigation Controller
Weapons Systems
Defense Systems
Attacker

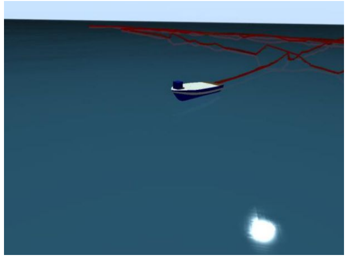

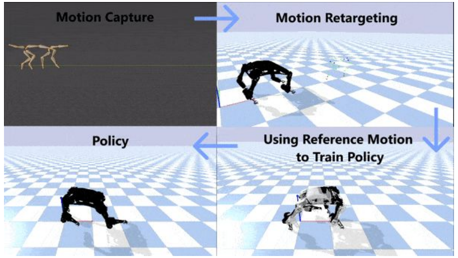

Motion Capture
Motion Retargeting
Policy
Using Reference Motion
to Train Policy

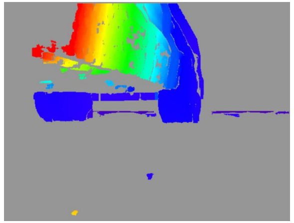

RTX 3080

Part 1: Tiny Robot Learning

Problem Motivation

Boston Dynamics Spot

Quadruped Robot

Problem Motivation

Unitree Go2 and B1 at iROS 2023

Methods for Locomotion

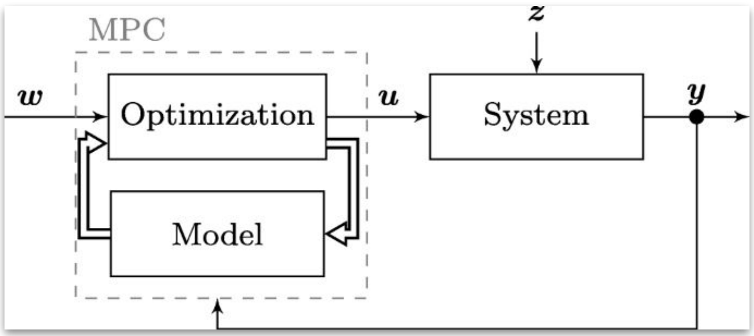

Model Predictive Control (MPC)

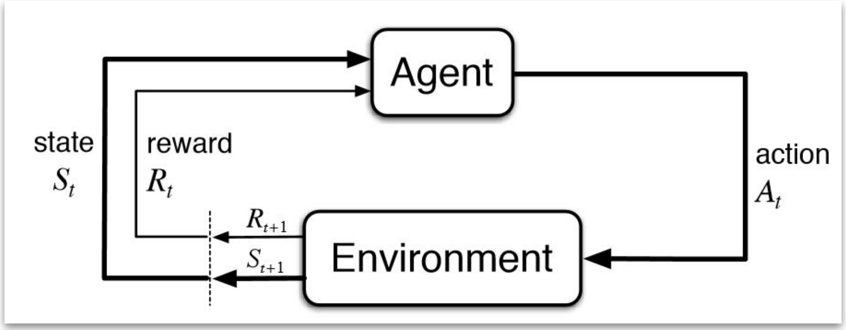

Reinforcement Learning (RL)

Methods for Locomotion

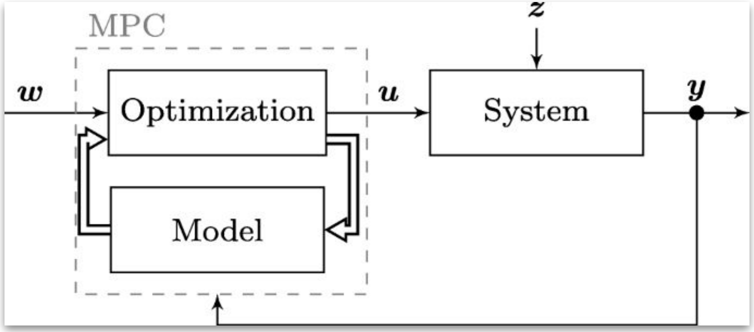

Model Predictive Control (MPC)

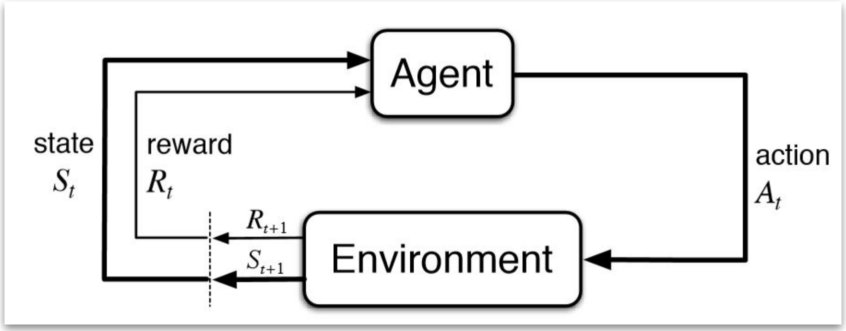

Reinforcement Learning (RL)

What are some advantages and disadvantages of each method?

RL & Natural Gaits

RL & Natural Gaits

What causes this type of RL behavior?

Imitation Learning

Learning Agile Robotic Locomotion Skills by Imitating Animals

Xue Bin Peng*†, Erwin Coumans*, Tingnan Zhang*, Tsang-Wei Edward Lee*, Jie Tan*, Sergey Levine*†
*Google Research, †University of California, Berkeley
Email: xbpeng@berkeley.edu, {erwincoumans,tingnan,tsangwei,jietan}@google.com, svlevine@eecs.berkeley.edu

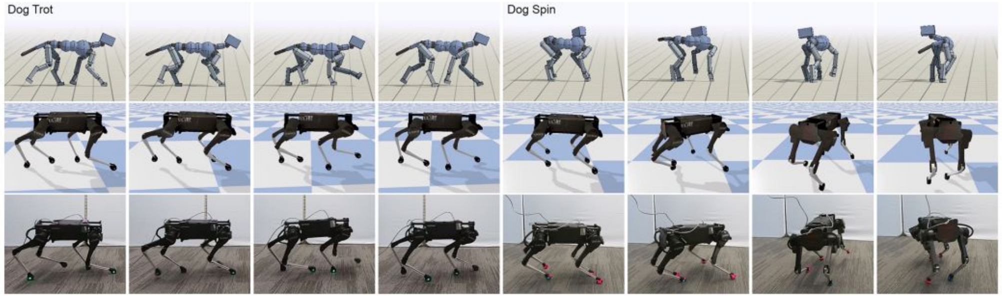

Fig. 1. Laikago robot performing locomotion skills learned by imitating motion data recorded from a real dog. **Top:** Motion capture data recorded from a dog. **Middle:** Simulated Laikago robot imitating reference motions. **Bottom:** Real Laikago robot imitating reference motions.

Abstract—**Reproducing the diverse and agile locomotion skills of animals has been a longstanding challenge in robotics. While manually-designed controllers have been able to emulate many**

designing control strategies often involves a lengthy development process, and requires substantial expertise of both the

https://xbpeng.github.io/projects/Robotic_Imitation/2020_Robotic_Imitation.pdf

Imitation Learning

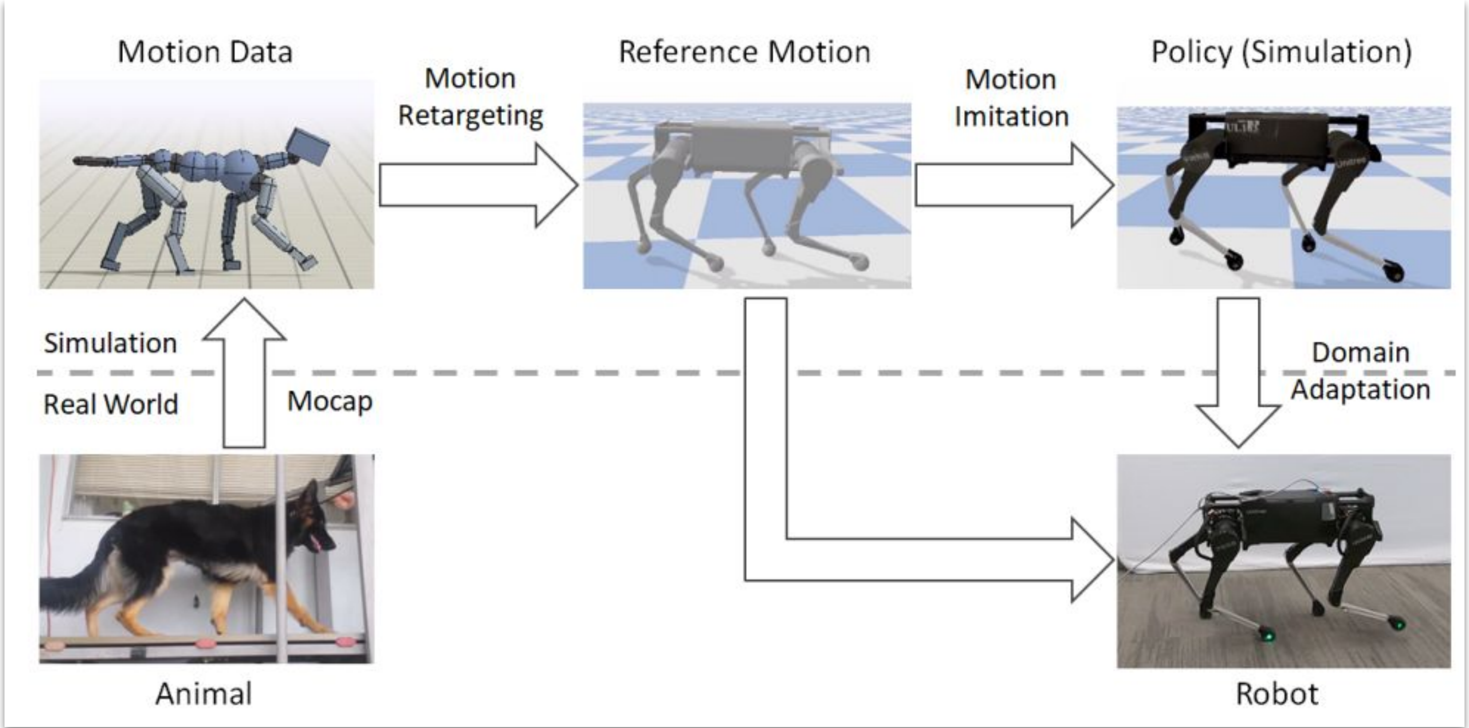

https://xbpeng.github.io/projects/Robotic_Imitation/2020_Robotic_Imitation.pdf

Imitation Learning & Natural Gaits

Quadruped Comparison

Unitree A1

Petoi Bittle

Can you spot some differences?

Quadruped Comparison: **Price**

Boston Dynamics Spot

Unitree A1

Unitree Go2

Petoi Bittle

Quadruped Comparison: **Price**

Boston Dynamics Spot

> $100k

Unitree A1

~ $10k

Unitree Go2

~ $2.5K

Petoi Bittle

$299

Quadruped Comparison: **Price**

33x

Ultra Low Cost

Boston Dynamics Spot

Unitree A1

Unitree Go2

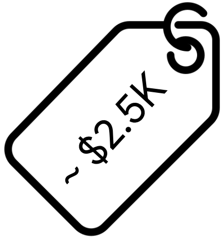

Petoi Bittle

Quadruped Comparison: **Size**

Unitree A1

1.64 ft x 0.98 ft x 1.31 ft

Petoi Bittle

7.8 in x 4.3 in x 4.3 in

2.5x

Quadruped Comparison: **Compute**

Unitree A1

Petoi Bittle

Standard Compute:	ARM Cortex-A72 2.5GHz	**125x**	Nyboard V1 ATMega328P 20MHz
Additional Compute:	NVIDIA TX2 1.3GHz	**1.3x**	Raspberry Pi Zero 2W 1GHz

Quadruped Full Comparison

	Unitree A1	Petoi Bittle	Ratio
Cost	$10,000 USD	$299 USD	33x
Weight	12 kg	.29 kg	41x
Dimensions	.5 x .3 x .4 m	.2 x .11 x .11 m	2.5x
Degrees of Freedom (DoF)	12 (Leg: 3)	8 (Leg: 2)	1.5x
Battery Capacity	25.2V 4200mAh	7.4V 1000mAh	3x
Motor Resolution	.022°	1°	45x
IMU	Yes	Yes	NA
Motor Feedback	Yes	No	NA
Foot Pressure Sensor	Yes	No	NA
LiDAR	Yes	No	NA
Computing	ARM Cortex-A72 2.5GHz	Nyboard V1 ATMega328P 20MHz	125x
Optional Additional Computing	NVIDIA TX2 1.3GHz	Raspberry Pi Zero 2W 1GHz	1.3x

Goal

Demonstrate how existing state-of-the-art **imitation learning pipelines** can be **modified** and augmented to support **ultra-low-cost, constrained robot** platforms

Challenges

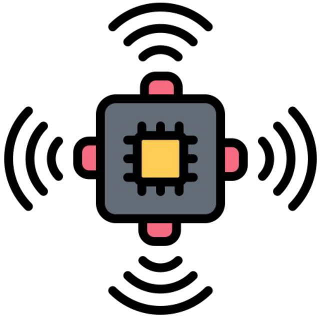

Observability

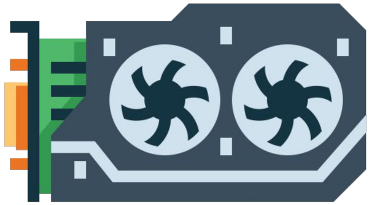

Computation

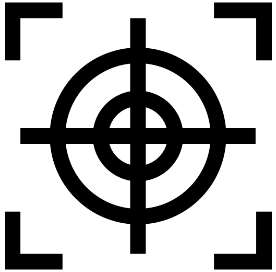

Controllability

Adjustments

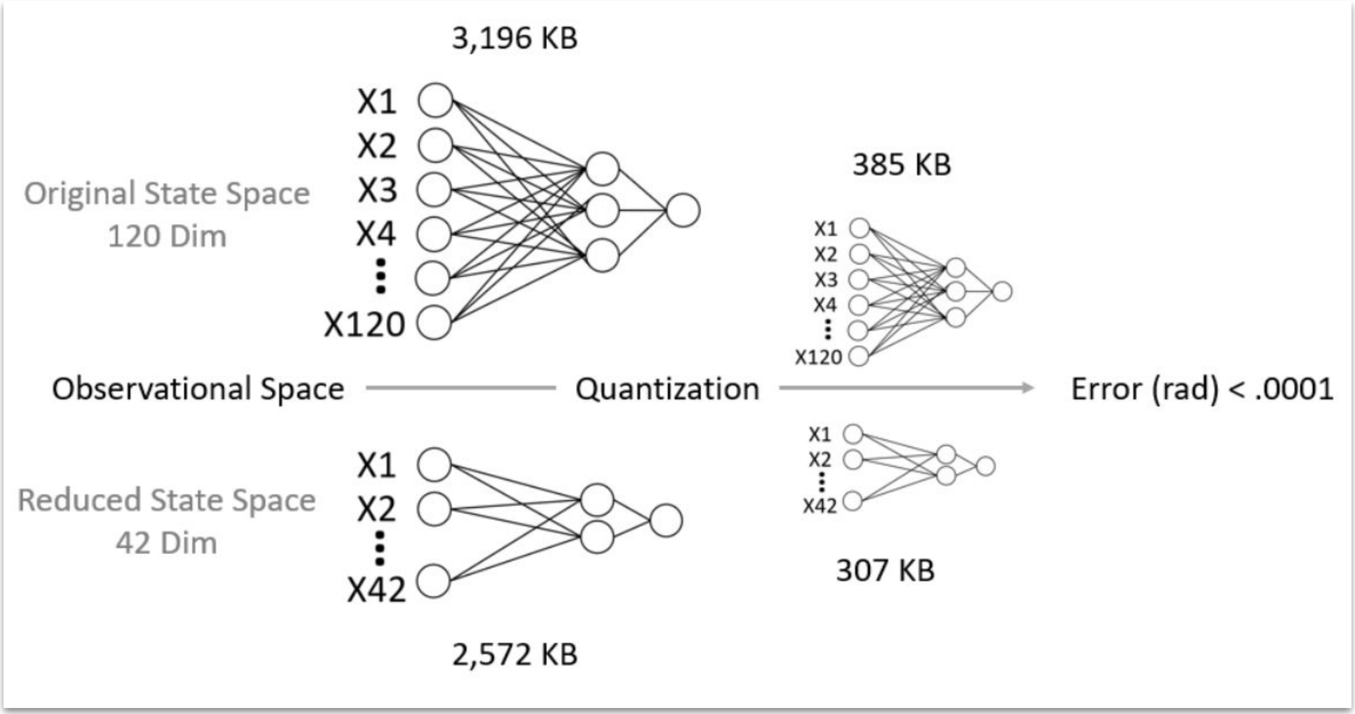

Adjustments

- IMU Data
- Motor Angles
- Previous Actions
- Future Reference Motion Frames

- Joint Angles

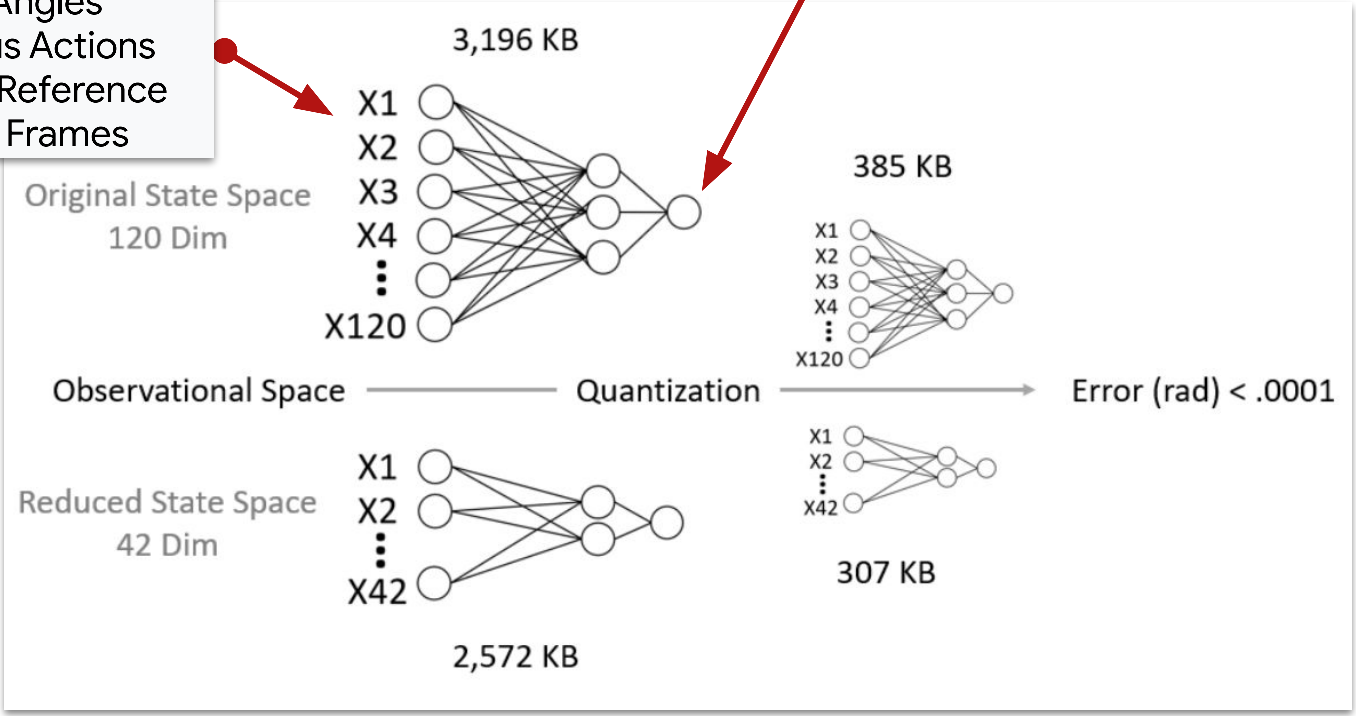

Adjustments

- IMU Data
- Motor Angles
- Previous Actions
- Future Reference Motion Frames

- Joint Angles

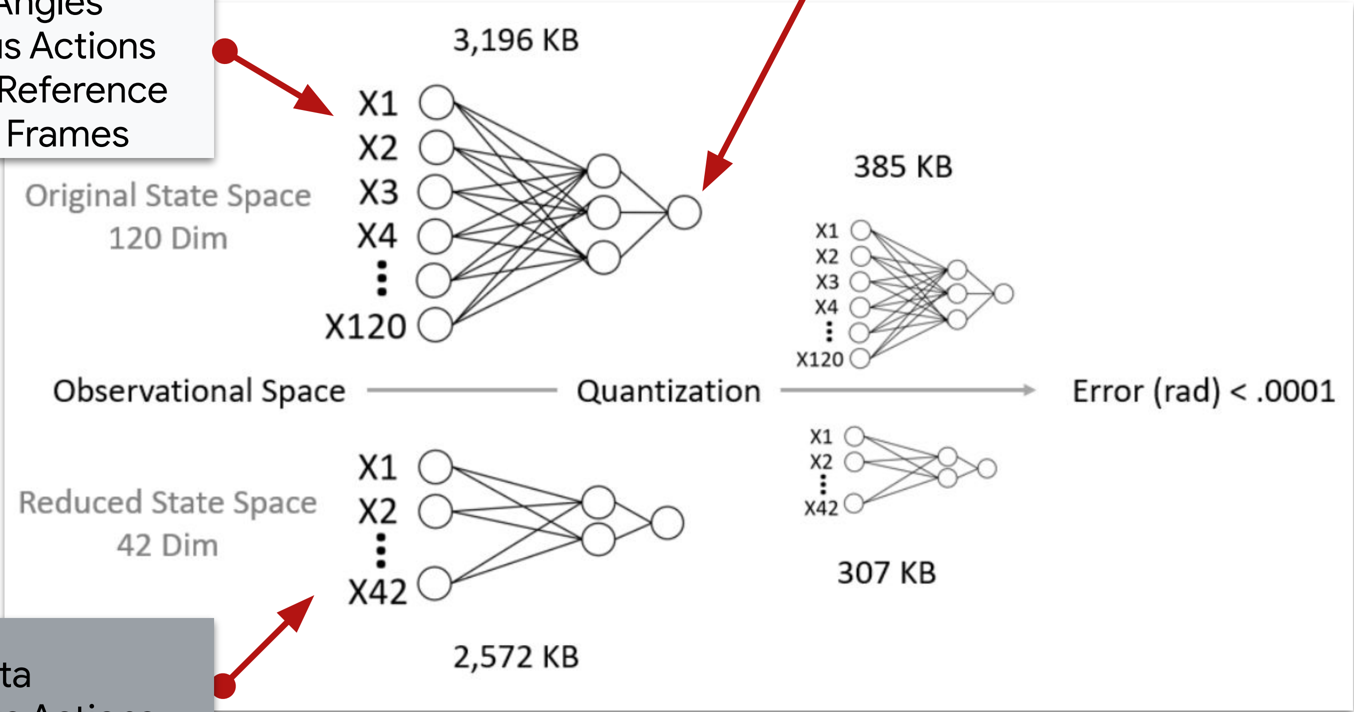

- IMU Data
- Previous Actions

Adjustments

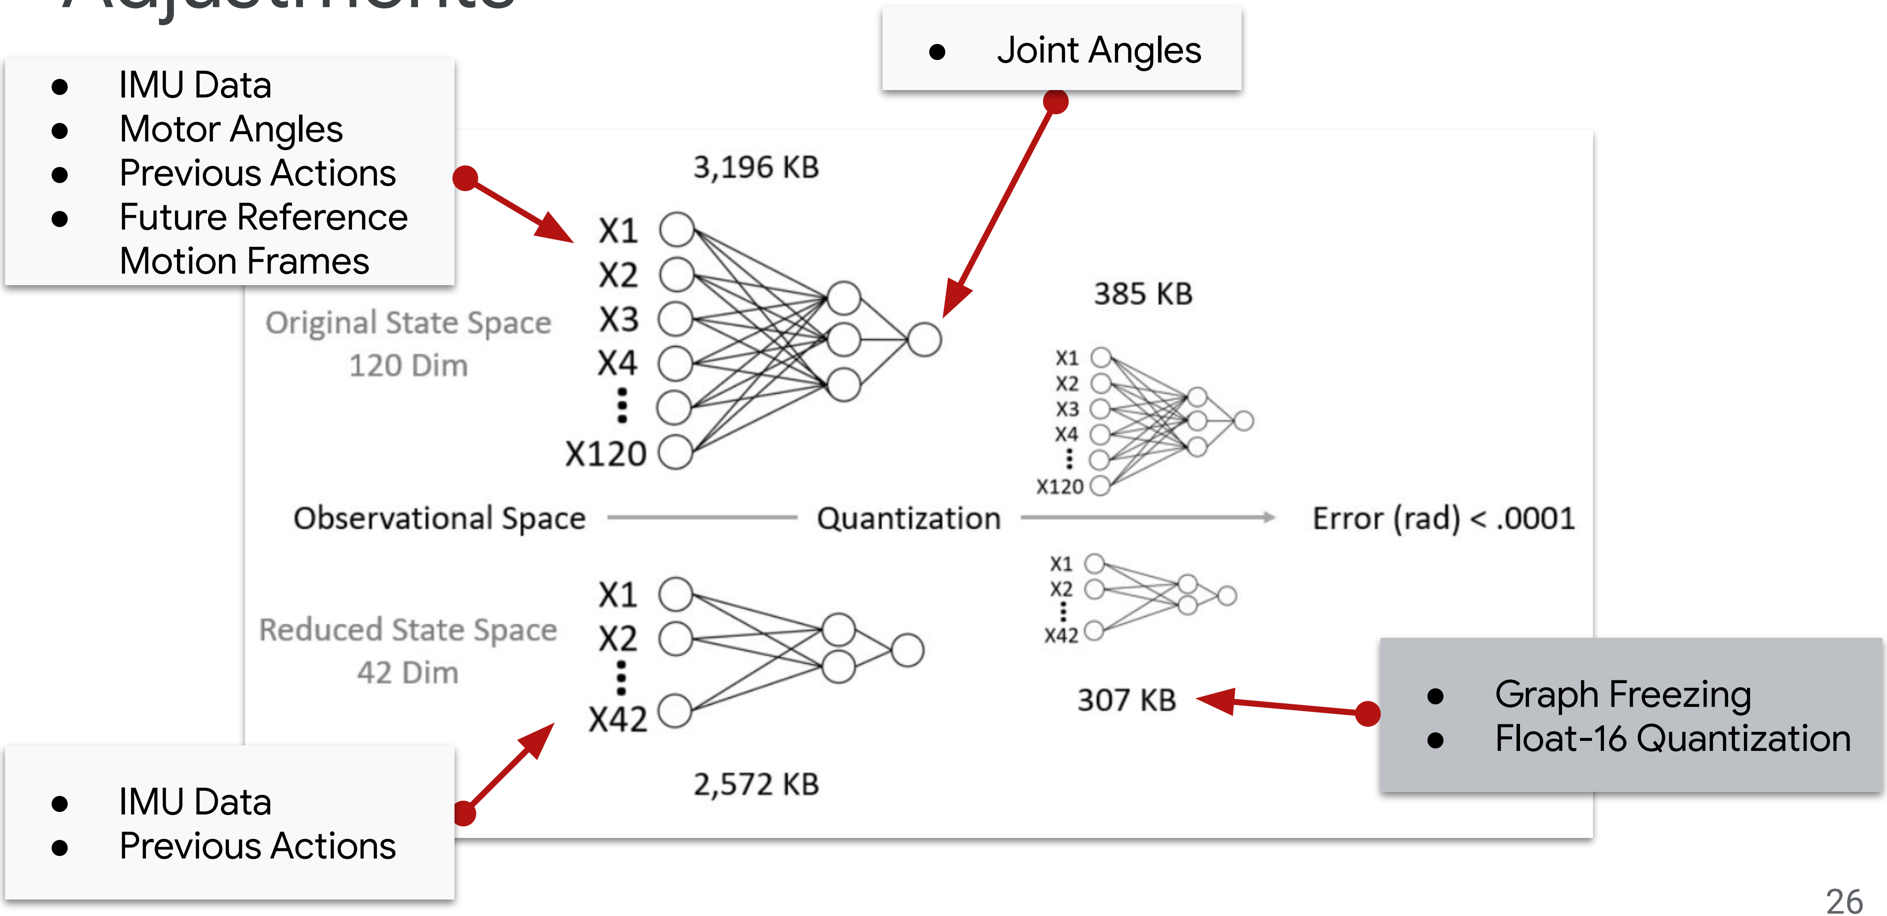

Adjustments

- IMU Data
- Motor Angles
- Previous Actions
- Future Reference Motion Frames

- Joint Angles

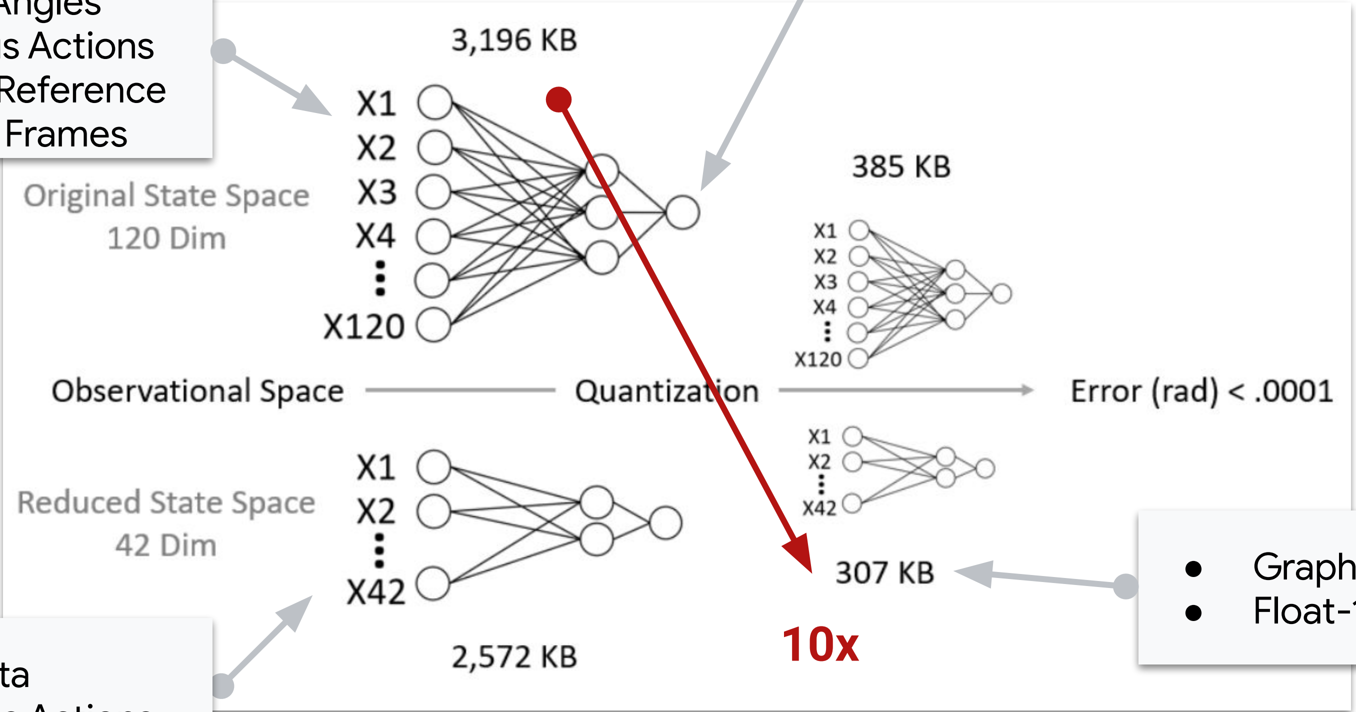

- Graph Freezing
- Float-16 Quantization

- IMU Data
- Previous Actions

Adjustments

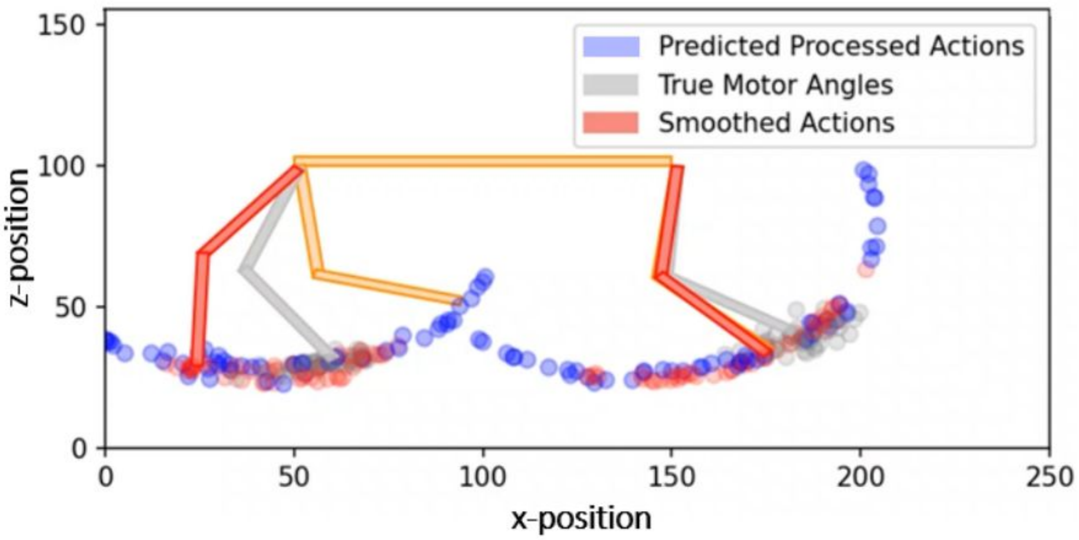

Tiny Motion Imitation Pipeline

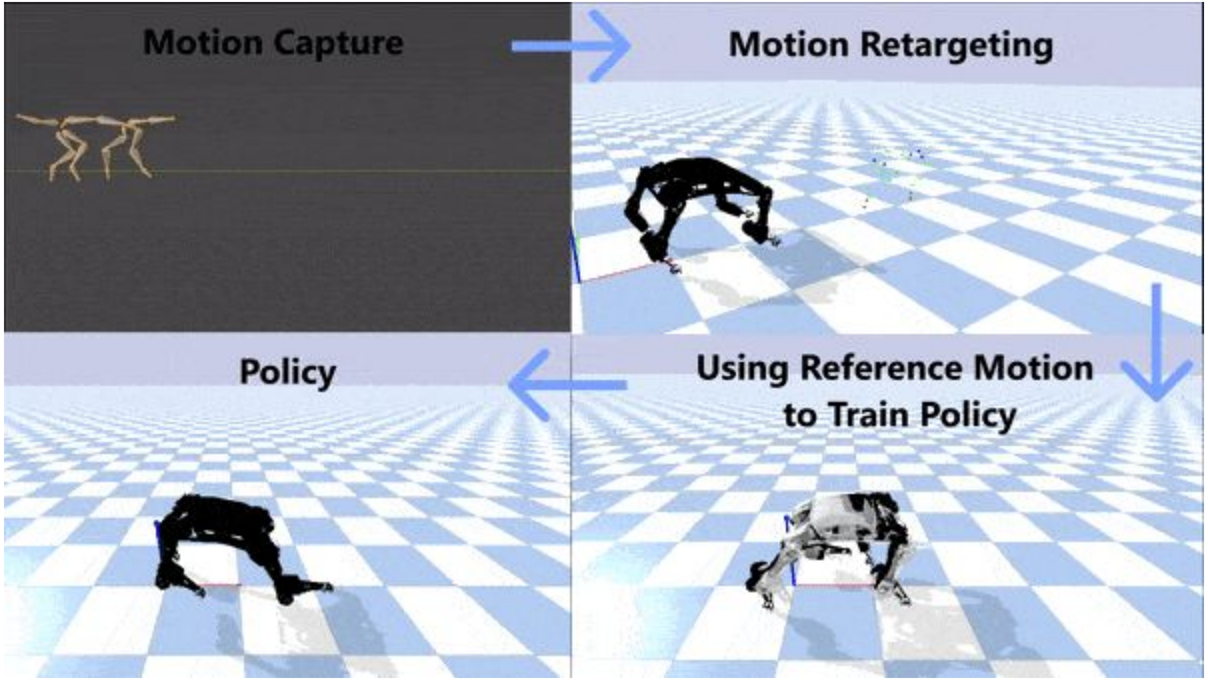

Tiny Motion Imitation Pipeline

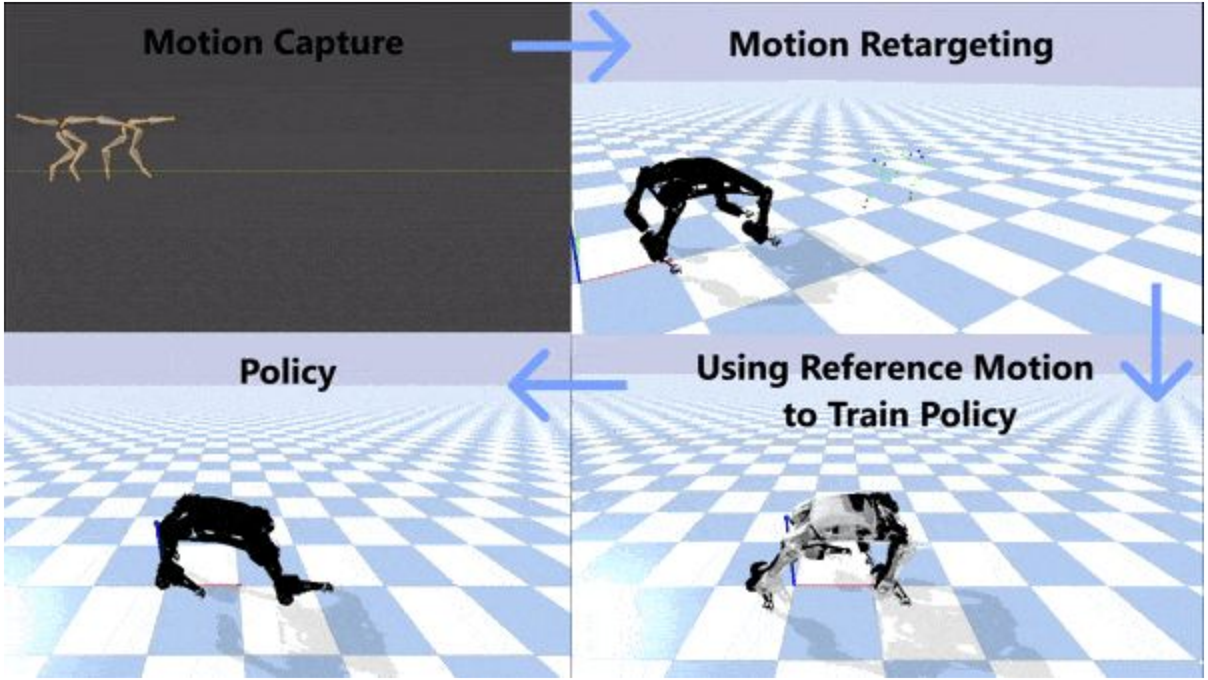

PPO1 for 60M Timesteps

Tiny Motion Imitation Pipeline

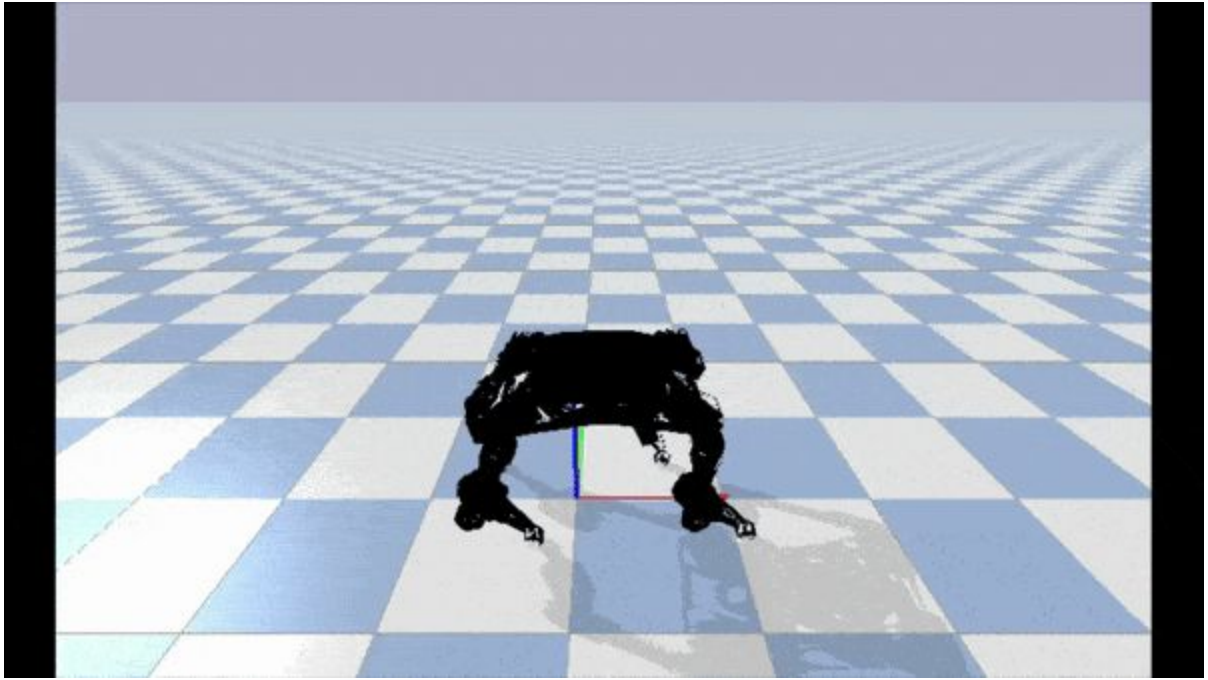

Quantized and Frozen Graph but without dimensionality reduction

Deployment

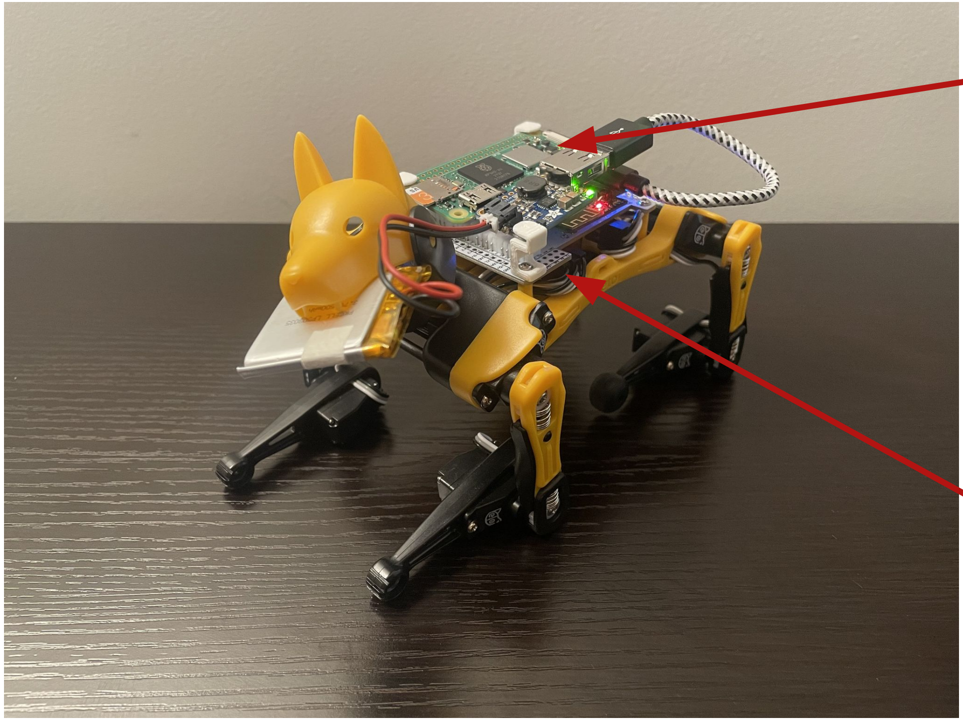

Raspberry Pi Zero 2W

Nyboard V1 ATMega328P

Video Demo

Current State-of-the-Art

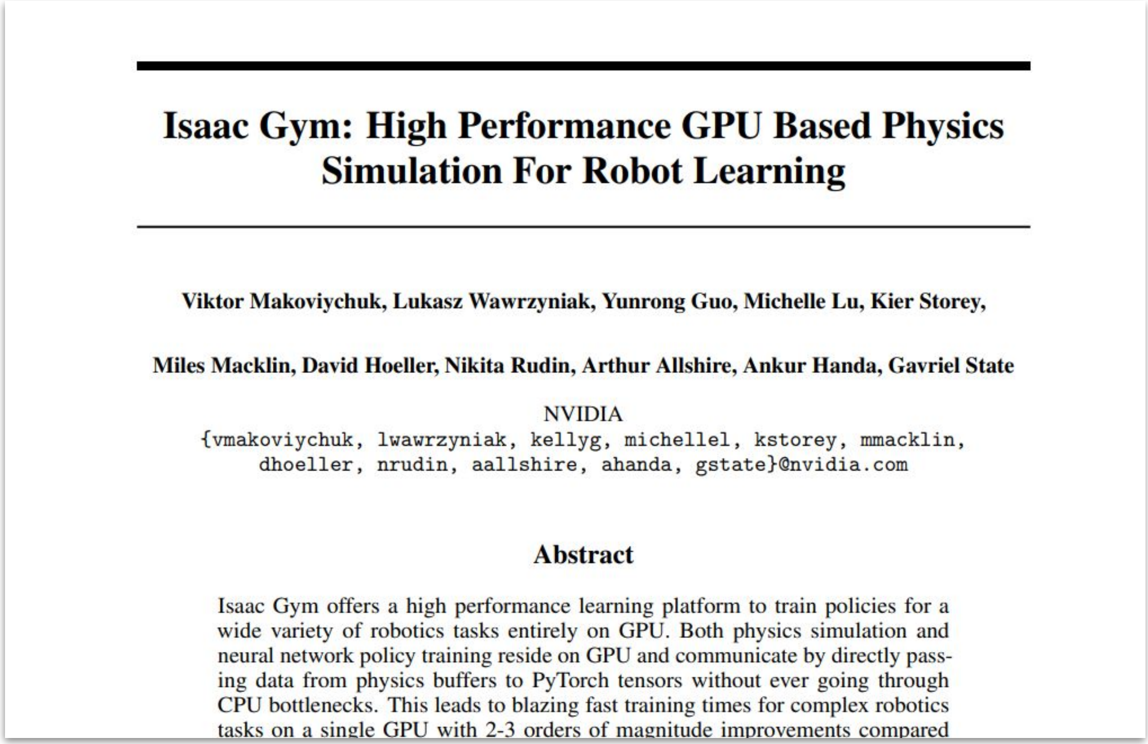

Isaac Gym: High Performance GPU Based Physics Simulation For Robot Learning

Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey,

Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, Gavriel State

NVIDIA
{vmakoviychuk, lwawrzyniak, kellyg, michellel, kstorey, mmacklin, dhoeller, nrudin, aallshire, ahanda, gstate}@nvidia.com

Abstract

Isaac Gym offers a high performance learning platform to train policies for a wide variety of robotics tasks entirely on GPU. Both physics simulation and neural network policy training reside on GPU and communicate by directly passing data from physics buffers to PyTorch tensors without ever going through CPU bottlenecks. This leads to blazing fast training times for complex robotics tasks on a single GPU with 2-3 orders of magnitude improvements compared

https://arxiv.org/pdf/2108.10470.pdf

Questions for Thought

- We rely heavily on simulation for training, but transferring these policies to the real world is still challenging. How big of a role do you think simulation can realistically play in training robots compared to real world training? What innovations might help close the sim2real gap?
- The locomotion skills in this paper are learned through imitation and reinforcement learning. Do you think techniques from general artificial intelligence like reasoning could complement these methods for robot training?
- Can you think of any other potential applications for capable yet affordable robots? What opportunities and risks might this create?

Part 2: RobotPerf

Our Team


Víctor Mayoral-Vilches

Jason Jabbour


Yu-Shun Hsiao


Zishen Wan


Alejandra Martínez-Fariña


Martino Crespo-Álvarez


Matthew Stewart


Juan Manuel Reina-Muñoz


Prateek Nagras


Gaurav Vikhe


Mohammad Bakhshalipour


Martin Pinzger


Stefan Rass


Smruti Panigrahi


Giulio Corradi


Niladri Roy


Phillip B. Gibbons


Sabrina M. Neuman


Brian Plancher


Vijay Janapa Reddi

ACCELERATION ROBOTICS

HARVARD UNIVERSITY

Georgia Tech

Carnegie Mellon University

Ford

intel

JKU JOHANNES KEPLER UNIVERSITÄT LINZ

BARNARD

UNIVERSITÄT KLAGENFURT

BOSTON UNIVERSITY

AMD

Introduction

Motivation

Robotic Applications

Real-Time Systems

Moore's Law & Dennard Scaling

Heterogeneous Hardware

Overview

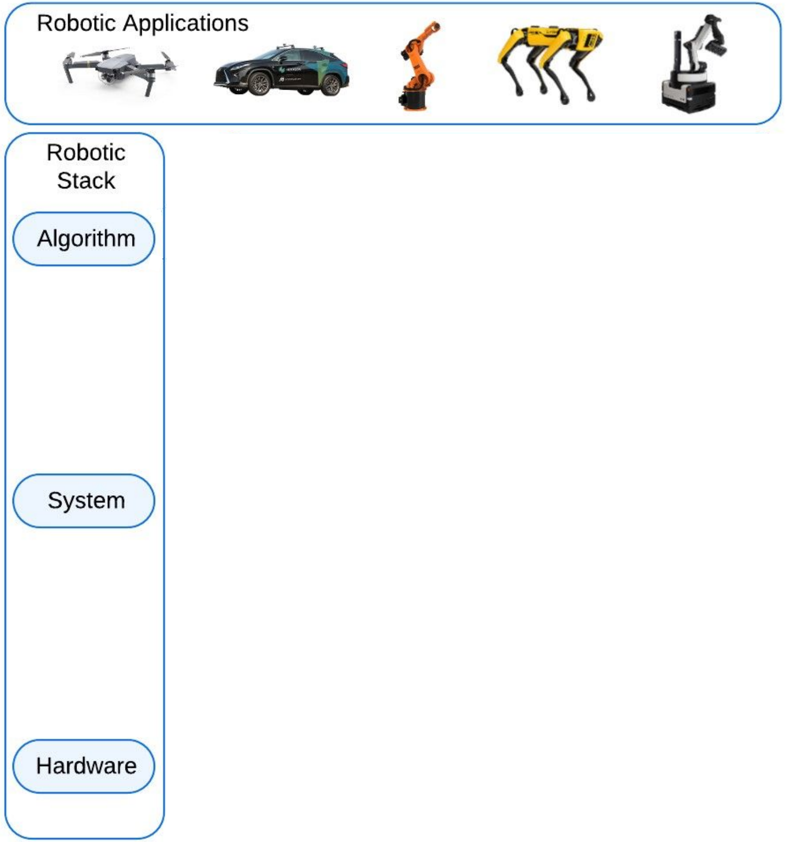
Robotic Applications
Robotic Stack
Algorithm
System
Hardware

Overview

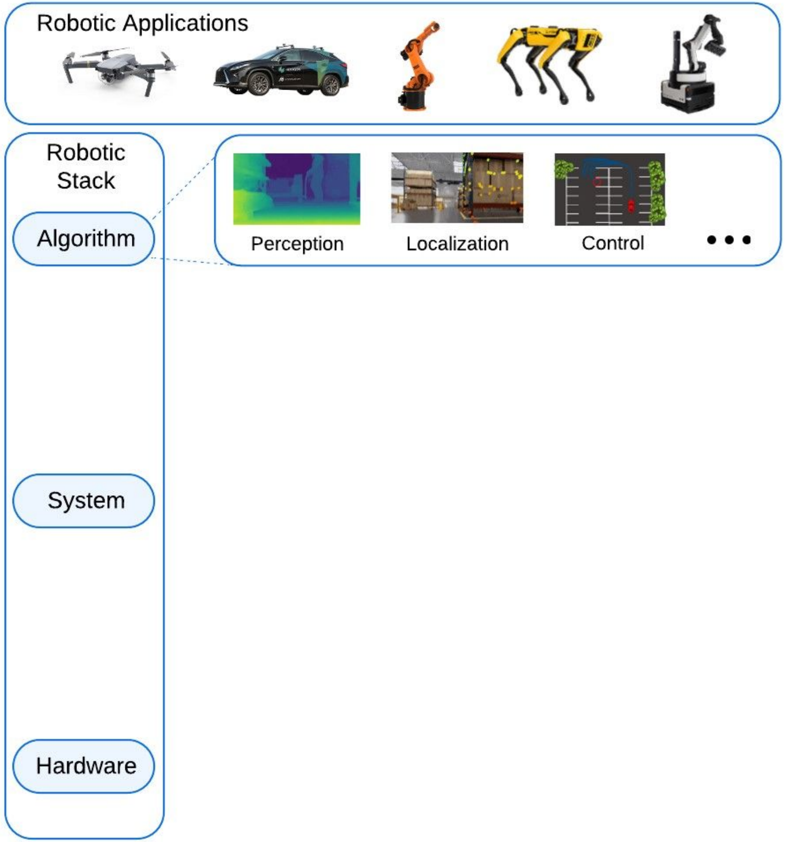

Robotic Applications
Robotic Stack
Algorithm
Perception
Localization
Control
• • •
System
Hardware

Overview

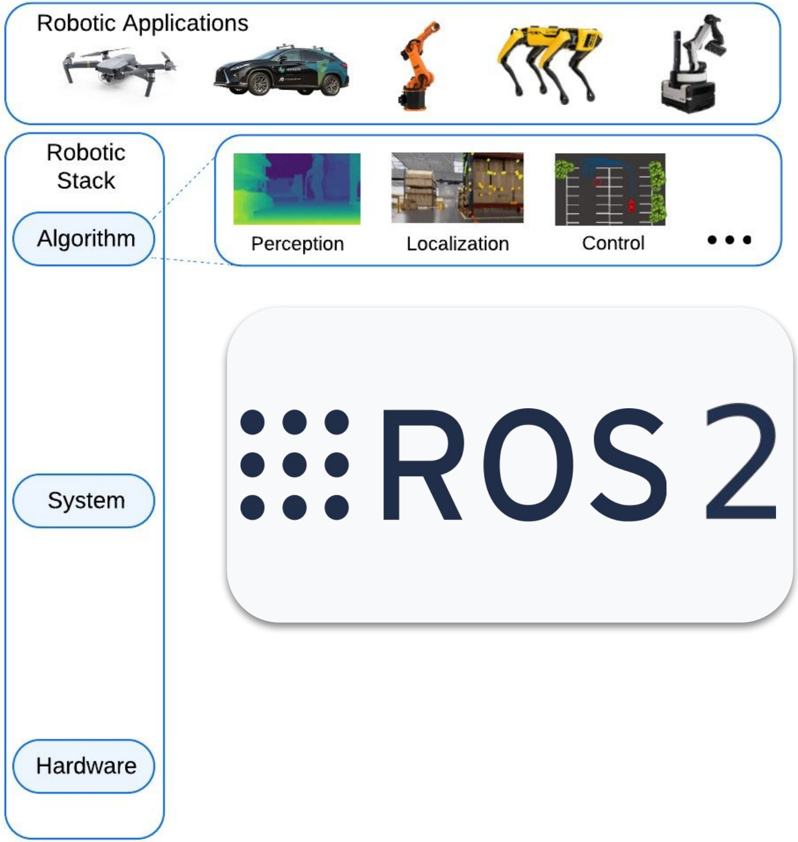
Robotic Applications
Robotic Stack
Algorithm
Perception
Localization
Control
• • •
System
Hardware
ROS 2

ROS 2

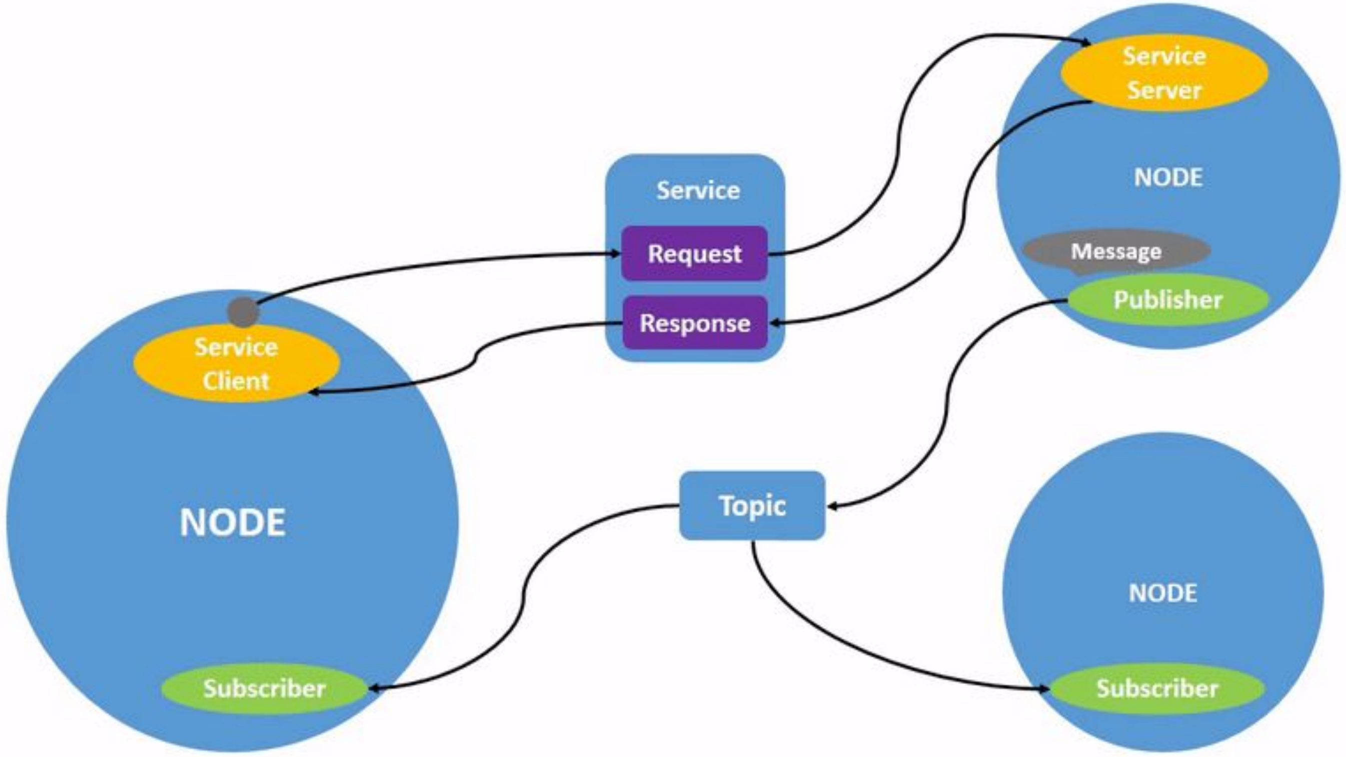
Service
Server
NODE
Message
Publisher
Service
Request
Response
Service
Client
NODE
Topic
NODE
Subscriber
Subscriber

Overview

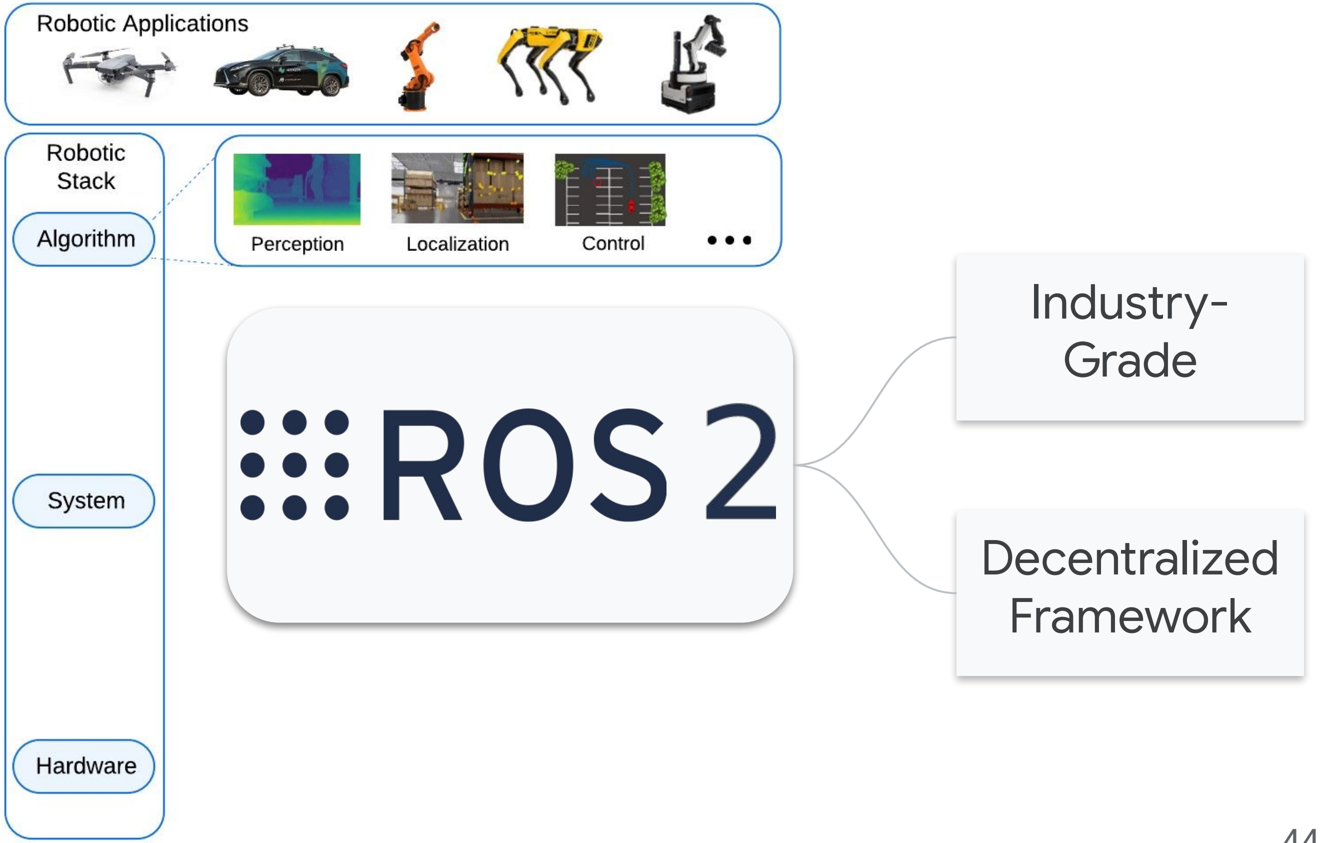
Robotic Applications
Robotic Stack
Algorithm
Perception
Localization
Control
• • •
System
Hardware
ROS 2
Industry-Grade
Decentralized Framework

Background

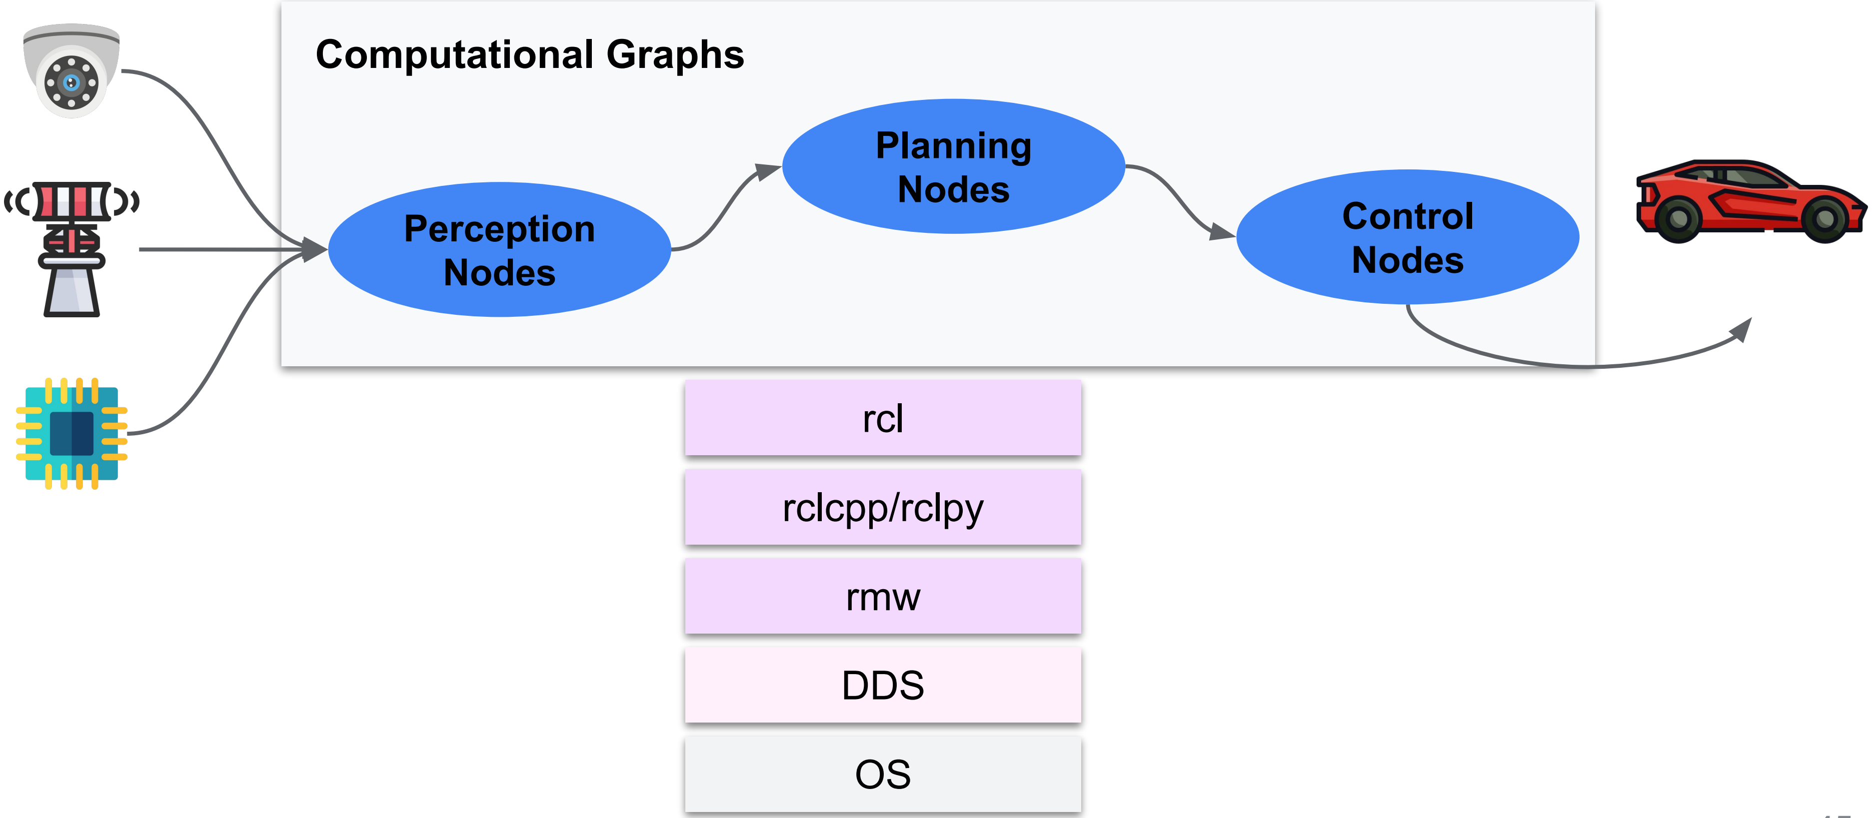
Computational Graphs
Perception Nodes
Planning Nodes
Control Nodes
rcl
rclcpp/rclpy
rmw
DDS
OS

Overview

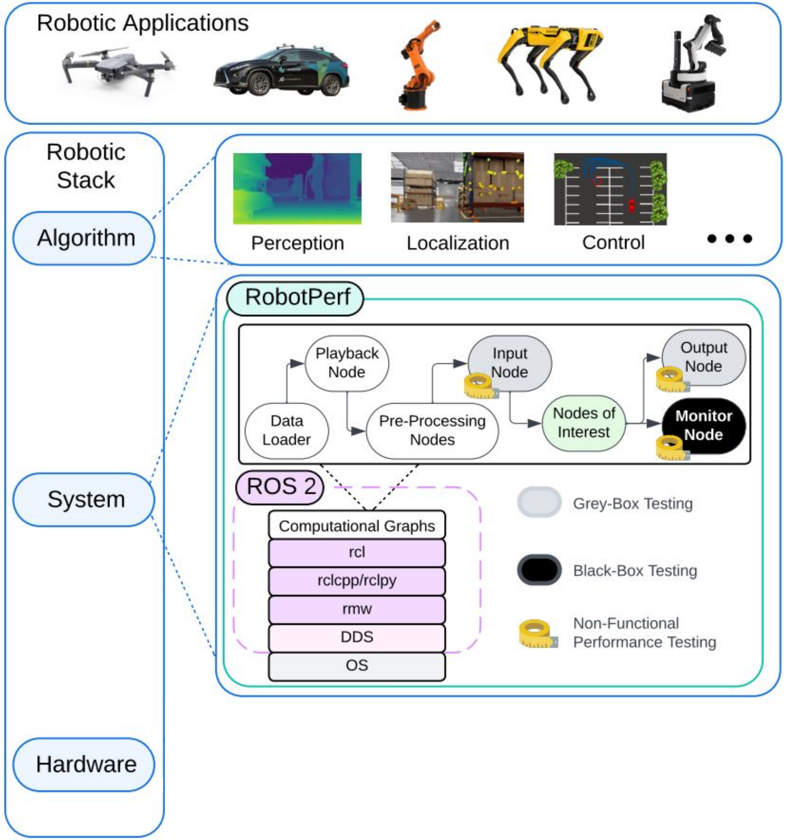
Robotic Applications
Robotic Stack
Algorithm
System
Hardware
Perception
Localization
Control
• • •
RobotPerf
Playback Node
Input Node
Output Node
Data Loader
Pre-Processing Nodes
Nodes of Interest
Monitor Node
ROS 2
Computational Graphs
rcl
rclcpp/rclpy
rmw
DDS
OS
Grey-Box Testing
Black-Box Testing
Non-Functional Performance Testing

Overview

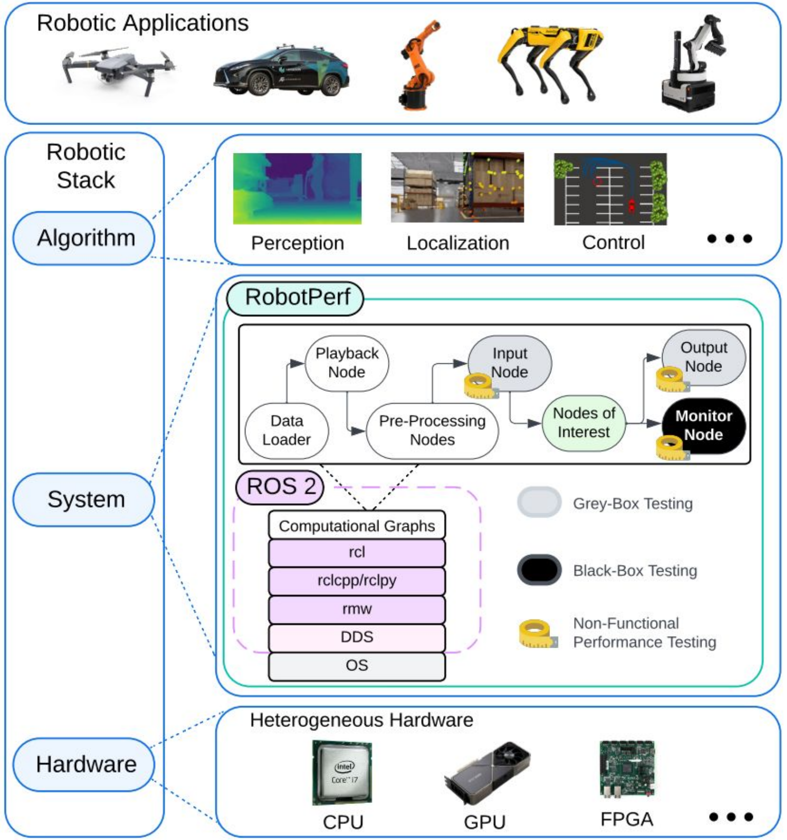

Robotic Applications
Robotic Stack
Algorithm
Perception
Localization
Control
RobotPerf
Playback Node
Input Node
Output Node
Data Loader
Pre-Processing Nodes
Nodes of Interest
Monitor Node
ROS 2
System
Computational Graphs
rcl
rclcpp/rclpy
rmw
DDS
OS
Grey-Box Testing
Black-Box Testing
Non-Functional Performance Testing
Hardware
Heterogeneous Hardware
CPU
GPU
FPGA

Related Work

	Characteristics						
	Real-time Performance Metrics	Spans Multiple Pipeline Categories	Evaluation on Heterogeneous Hardware	Integration with ROS/ROS 2 Framework	Functional Performance Testing	Non-functional Performance Testing	Community Led
OMPL Benchmark [31]	✓	✗	✗	✗	✗	✓	✗
MotionBenchMaker [32]	✓	✗	✗	✗	✓	✓	✗
OpenCollBench [33]	✗	✗	✓	✗	✓	✗	✗
BARN [34]	✗	✗	✗	✓	✓	✗	✗
DynaBARN [35]	✓	✗	✗	✓	✓	✗	✗
MAVBench [25]	✓	✓	✓	✓	✓	✓	✗
Bench-MR [36]	✓	✗	✗	✗	✓	✗	✗
RTRBench [23]	✓	✓	✗	✗	✗	✓	✗
RobotPerf (ours)	✓	✓	✓	✓	✗	✓	✓

RobotPerf Principles

RobotPerf Principles

Non-Functional Performance Testing

RobotPerf
Benchmarks

Performance Testing Types

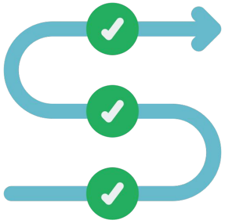

Functional

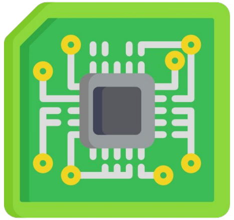

Non-Functional

RobotPerf Principles

Non-Functional Performance Testing

RobotPerf Benchmarks

Real-Time Metrics

Real-Time Metrics

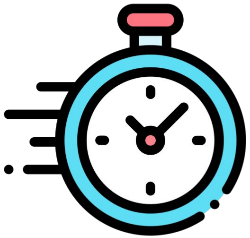

Latency

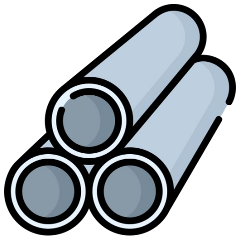

Throughput

Power

RobotPerf Principles

Non-Functional Performance Testing

Flexible Methodology

Real-Time Metrics

RobotPerf Benchmarks

Methodology Types

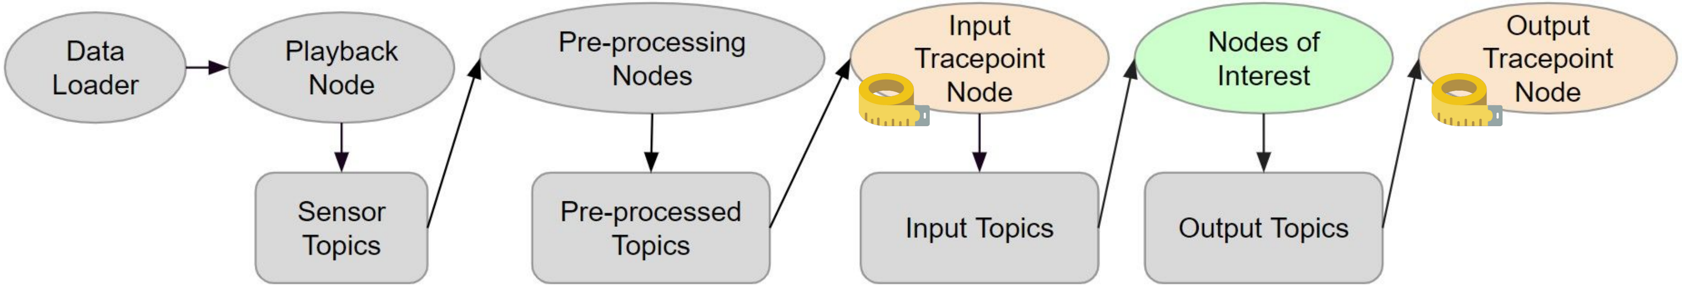

Grey Box Testing

Methodology Types

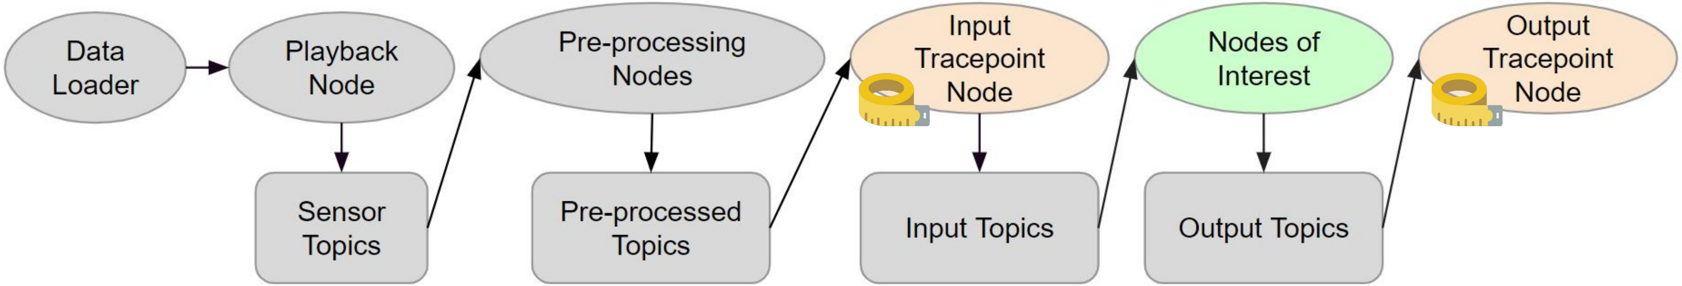

Grey Box Testing

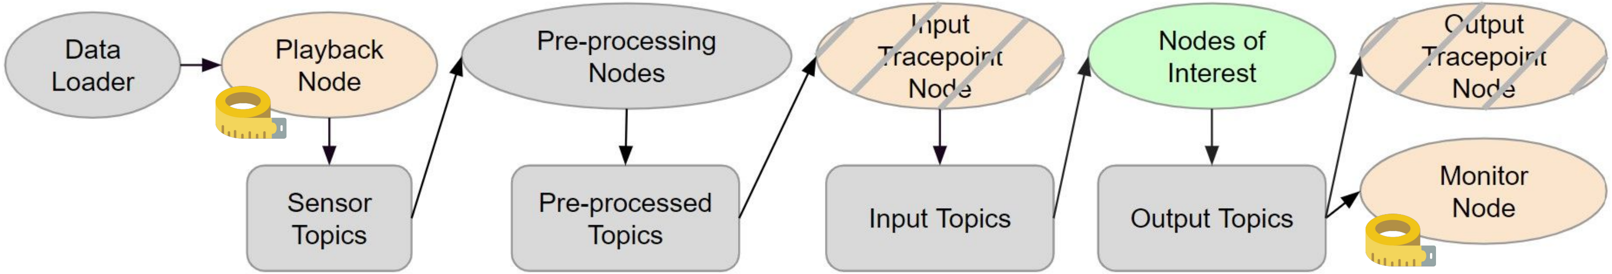

Black Box Testing

Methodology Types

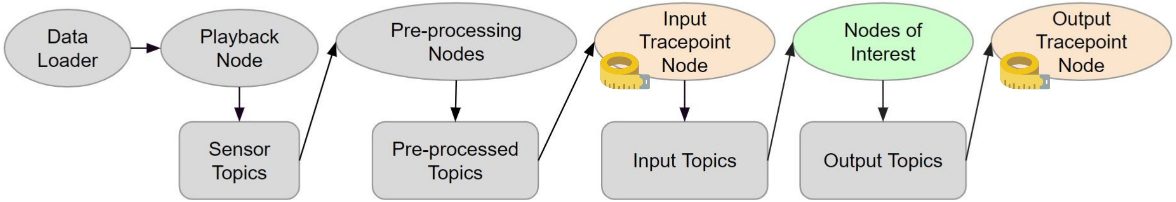

Grey Box Testing

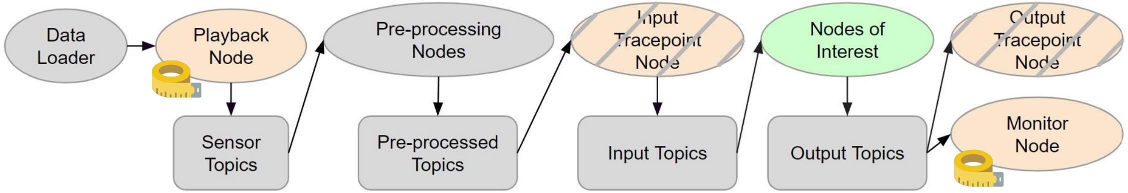

Black Box Testing

Tracing Granularity

Valid Tracer

Event Types

Code Modification

Standard ROS 2 APIs

Post-Processing

RobotPerf Principles

Non-Functional Performance Testing

Adaptability

RobotPerf Benchmarks

Flexible Methodology

Real-Time Metrics

Adaptability

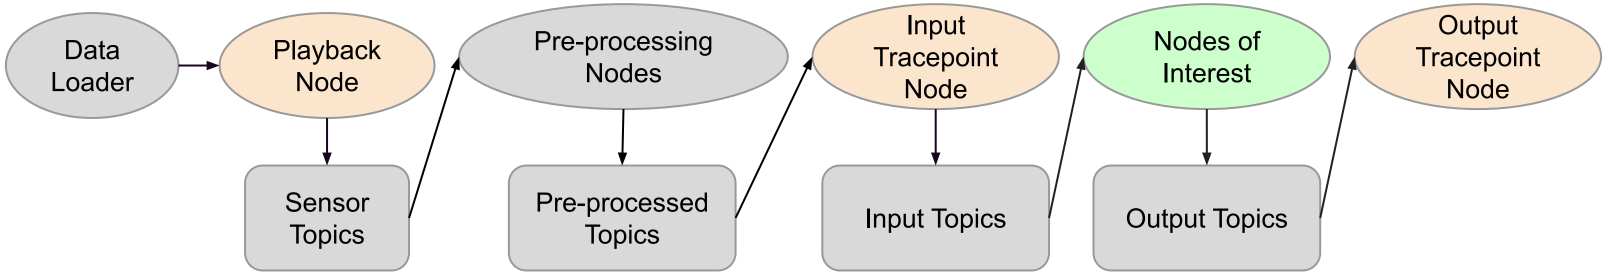

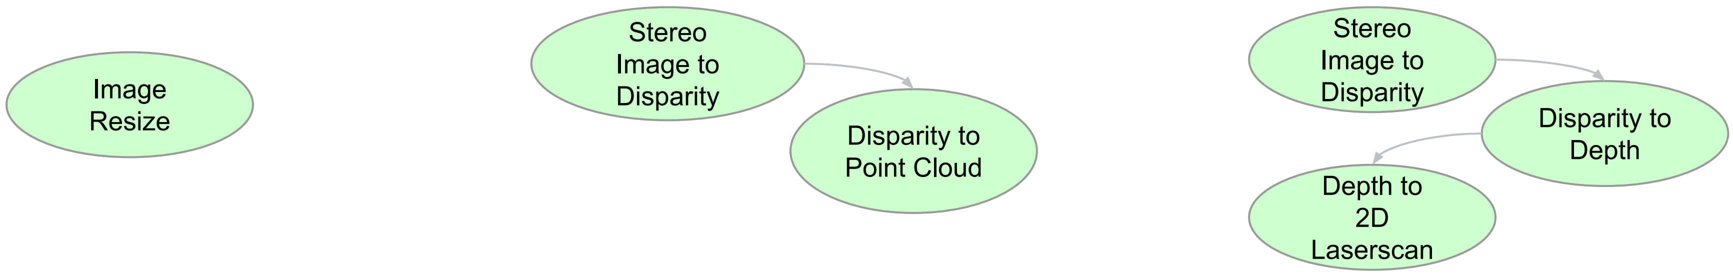

RobotPerf Principles

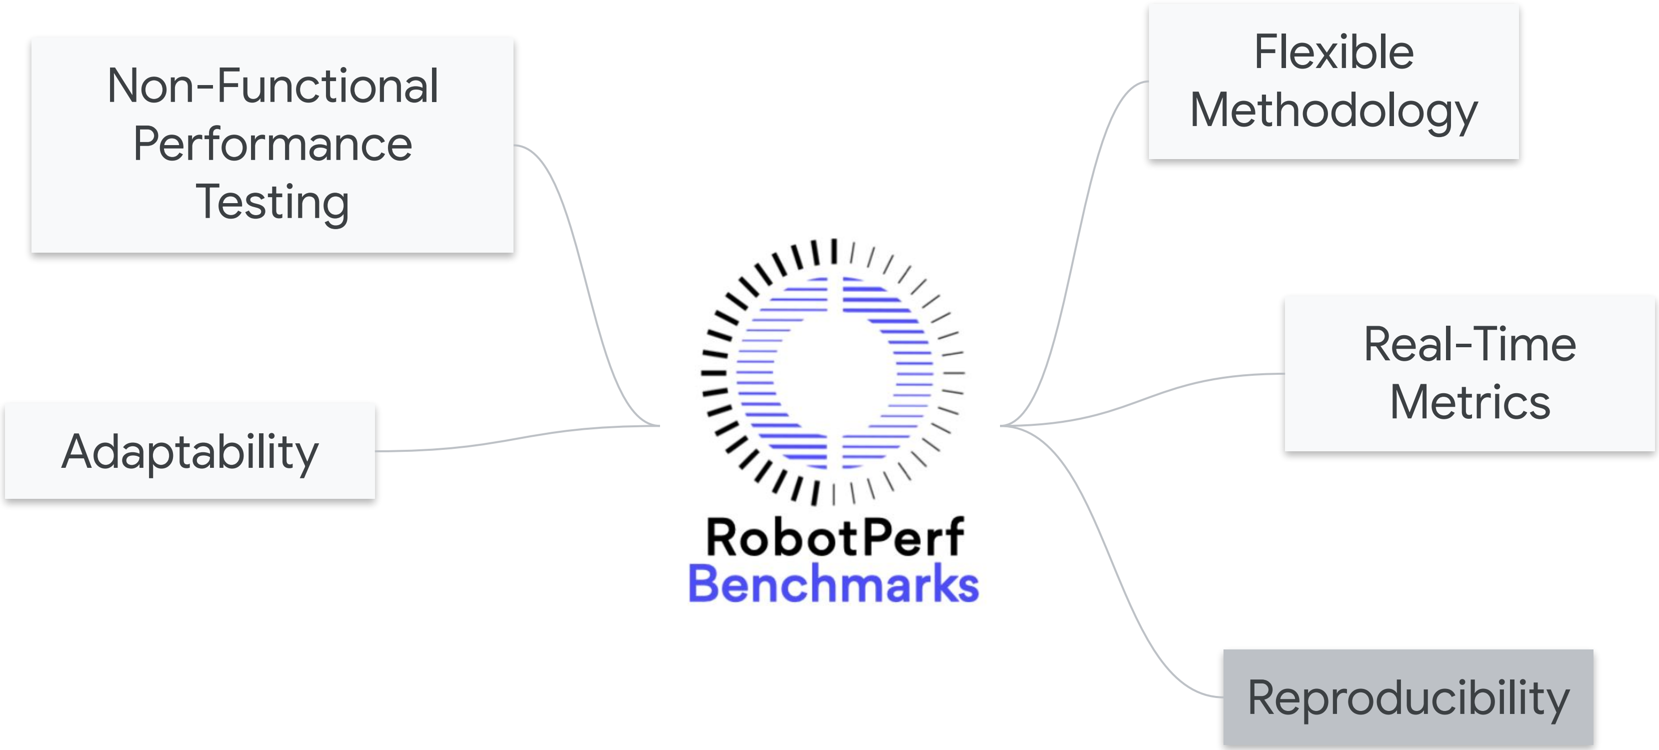

Reproducibility

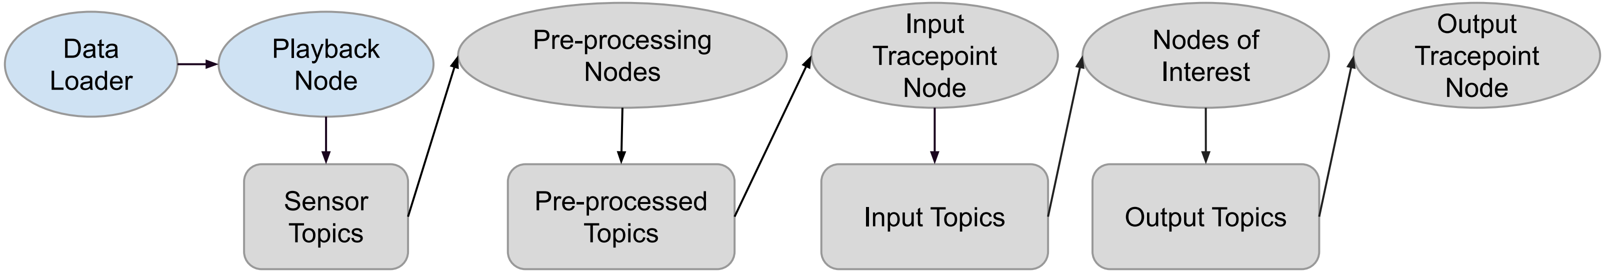

RobotPerf Principles

- Non-Functional Performance Testing
- Adaptability
- Portability
- Flexible Methodology
- Real-Time Metrics
- Reproducibility

RobotPerf
Benchmarks

RobotPerf Results

Benchmarks

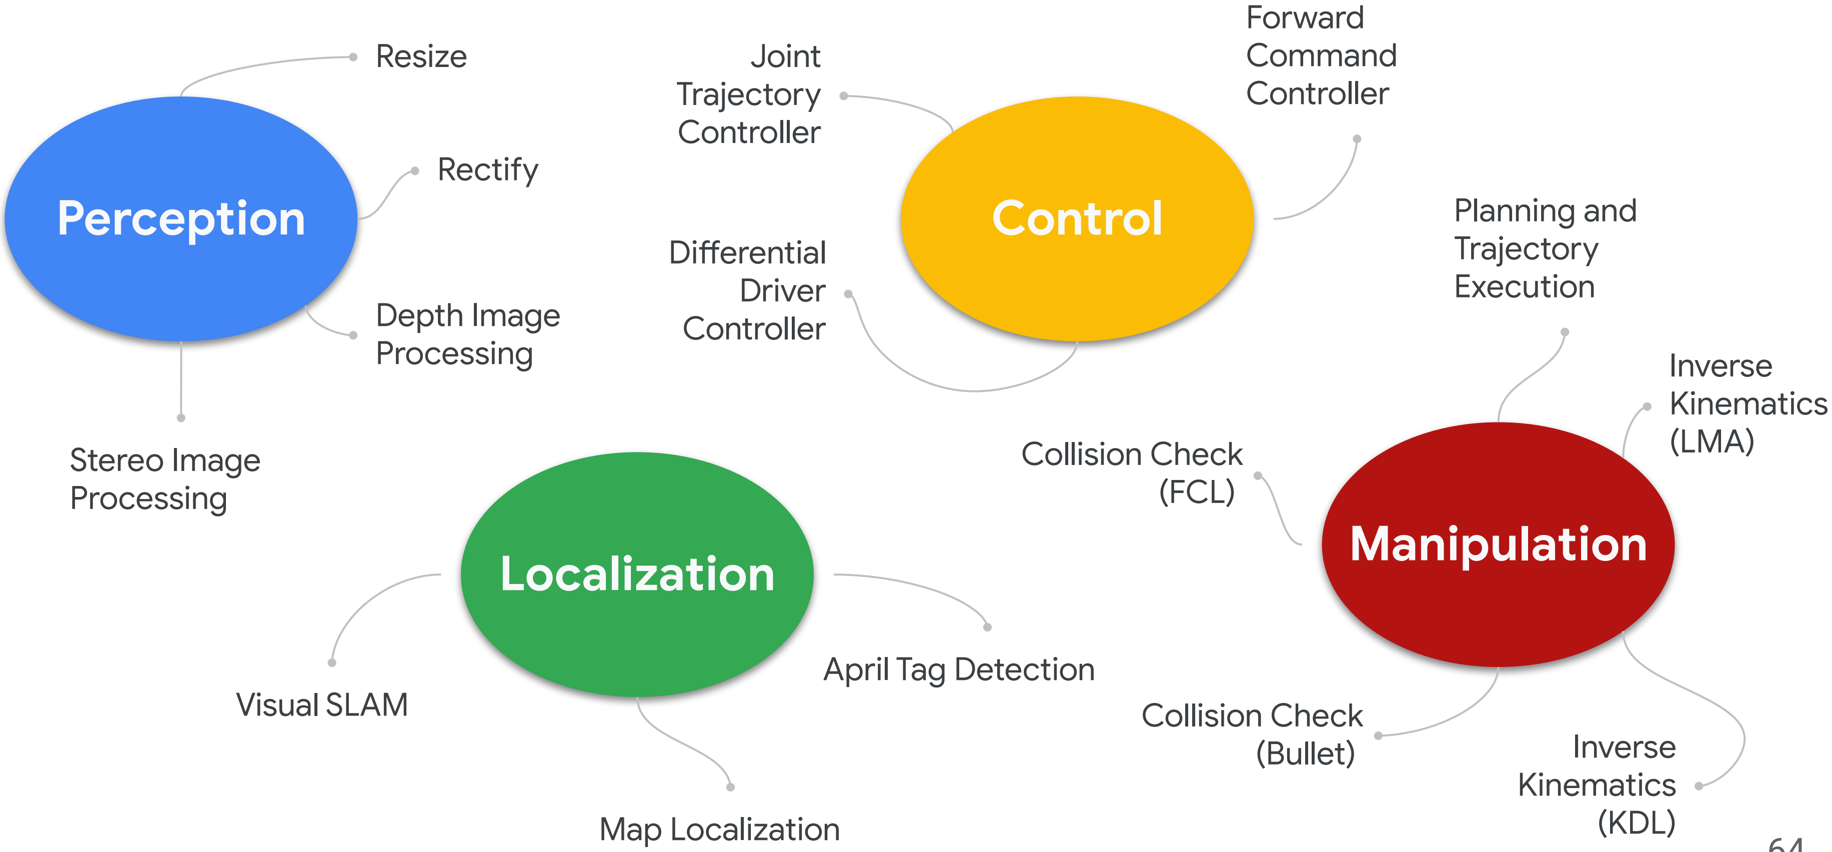
Perception
Resize
Rectify
Depth Image Processing
Stereo Image Processing
Control
Joint Trajectory Controller
Forward Command Controller
Differential Driver Controller
Localization
Visual SLAM
April Tag Detection
Map Localization
Manipulation
Planning and Trajectory Execution
Inverse Kinematics (LMA)
Collision Check (FCL)
Collision Check (Bullet)
Inverse Kinematics (KDL)

Takeaways

Quantitative Approach to Hardware Selection

Representative Assessment of Heterogeneous Hardware

Assessment of Acceleration Benefits

Hardware Selection

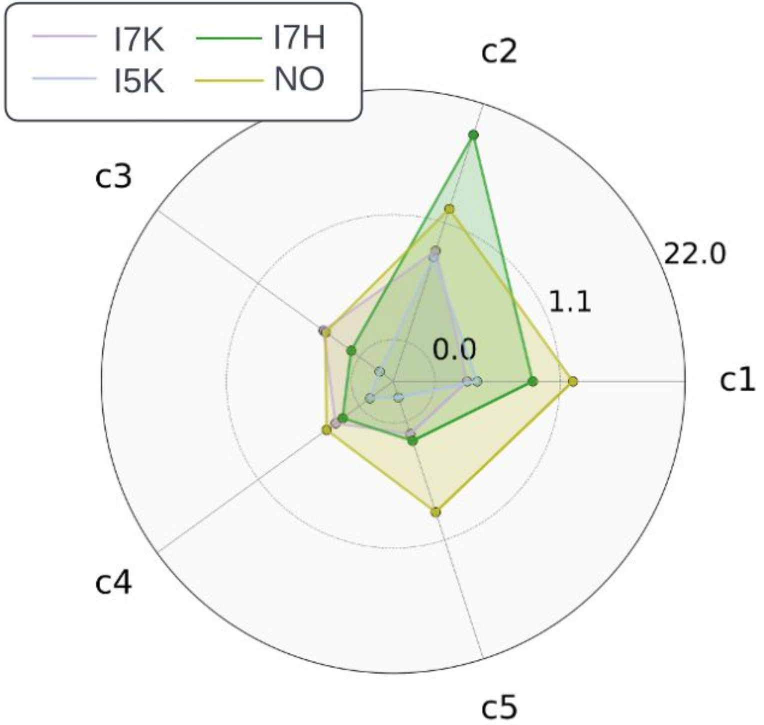

Representative Assessment

General-Purpose Hardware

- [I5U] (15W) Intel i5-8250U
- [AR] (65W) AMD Ryzen 5 PRO 4650G
- [I7K] (95W) Intel i7-8700K
- [I7H] (125W) Intel i7-12700H
- [I5K] (125W) Intel i5-13600K
- [I9K] (125W) Intel i9-12900KF

Heterogeneous Hardware

- [NN] (5W) NVIDIA Jetson Nano
- [QR] (5W) Qualcomm RB5 Robotics Kit
- [JX] (30W) Jetson AGX Xavier
- [NO] (60W) NVIDIA AGX Orin Dev. Kit
- [I7N] (295W) Intel i7-12700H + NVIDIA GeForce RTX 3060

Reconfigurable Hardware

- [KK] (15W) Kria KR260
- [KV] (36W) Kria KV260

Accelerator Hardware

- [KR] (35W) Kria KR260 (ROBOTCORE® Perception)
- [NR] (80W) NVIDIA AGX Orin Dev. Kit (ROBOTCORE® Perception)
- [I7T] (145W) Intel i7-12700H (ROBOTCORE® Transforms)
- [I5R] (315W) Intel i5-13600K + NVIDIA GeForce RTX 3060 (ROBOTCORE® Perception)
- [I7R] (315W) Intel i7-12700H + NVIDIA GeForce RTX 3060 (ROBOTCORE® Perception)

Grey Box

Black Box

Representative Assessment

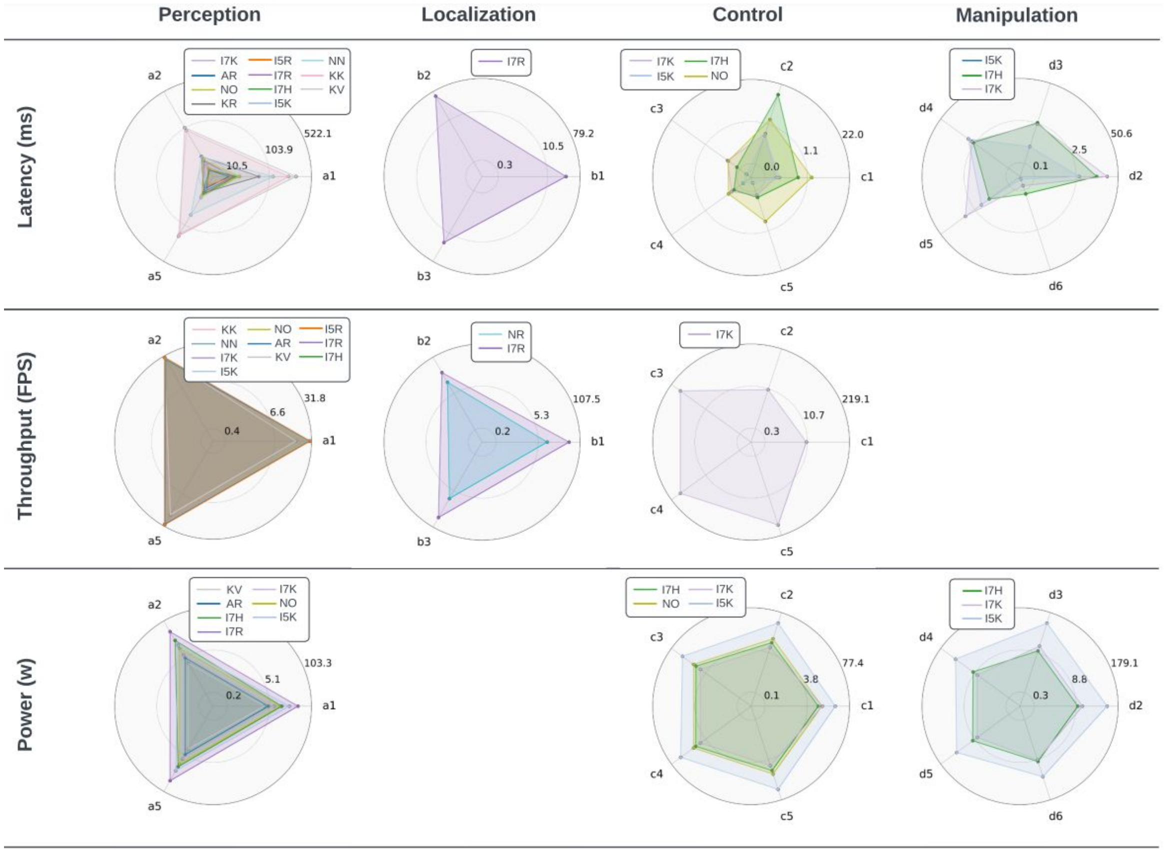

Acceleration Benefits

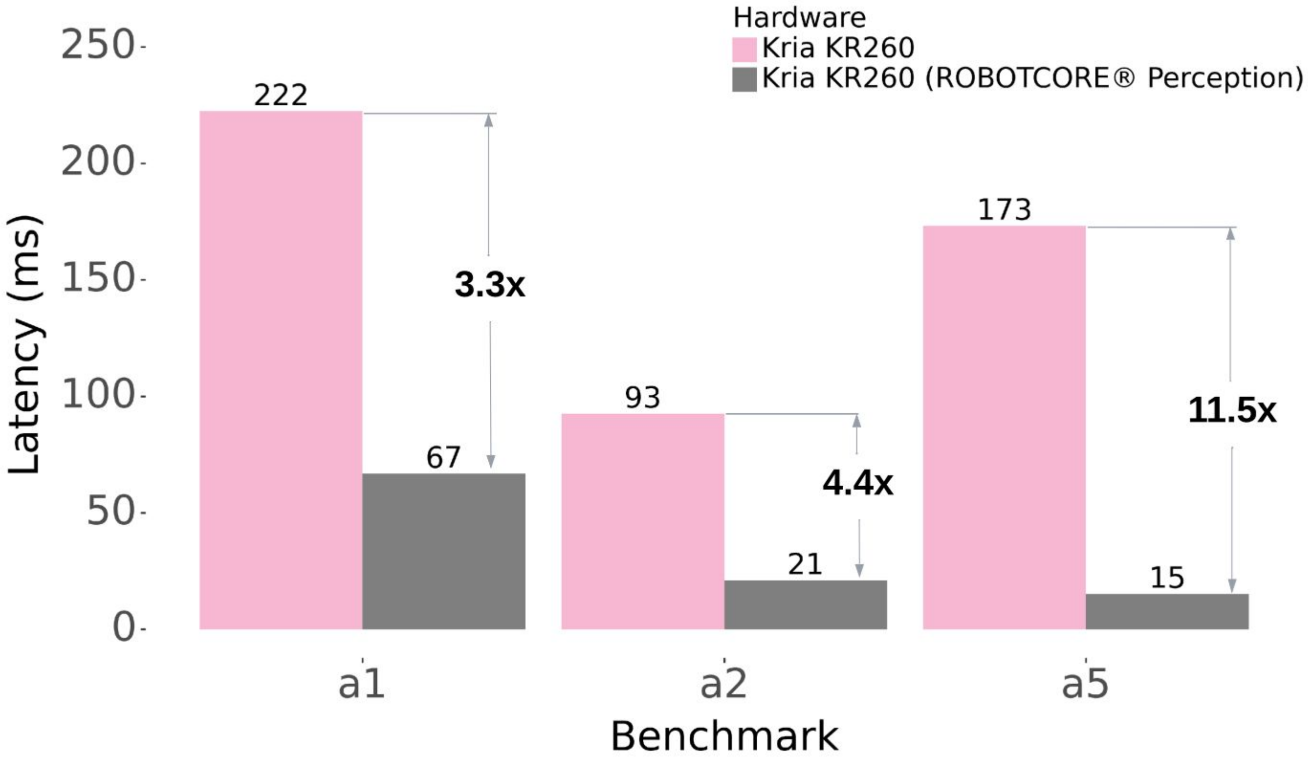

Thanks!

jasonjabbour@g.harvard.edu

Questions for Thought

- What's the importance of benchmarks? Can you think of some benchmarks that have helped move their field forward?
- Are there any specific robotic algorithms that should be incorporated into RobotPerf?
- Does RobotPerf miss or not take into consideration any aspects of the robotics pipeline that might be useful to study? What could be the next steps of RobotPerf?